

Identifying single electron charge sensor events using wavelet edge detection

J R Prance^{1,2}, B J Van Bael¹, C B Simmons¹, D E Savage¹,
M G Lagally¹, Mark Friesen¹, S N Coppersmith¹,
M A Eriksson¹

¹ University of Wisconsin-Madison, Wisconsin 53706, United States

² Department of Physics, Lancaster University, Lancaster, LA1 4YB, United Kingdom

Abstract. The operation of solid-state qubits often relies on single-shot readout using a nanoelectronic charge sensor, and the detection of events in a noisy sensor signal is crucial for high fidelity readout of such qubits. The most common detection scheme, comparing the signal to a threshold value, is accurate at low noise levels but is not robust to low-frequency noise and signal drift. We describe an alternative method for identifying charge sensor events using wavelet edge detection. The technique is convenient to use and we show that, with realistic signals and a single tunable parameter, wavelet detection can outperform thresholding and is significantly more tolerant to $1/f$ and low-frequency noise.

This is an author-created, un-copyedited version of an article accepted for publication/published in Nanotechnology. IOP Publishing Ltd is not responsible for any errors or omissions in this version of the manuscript or any version derived from it. The Version of Record is available online at doi:10.1088/0957-4484/26/21/215201.

1. Introduction

Charge sensors with single-electron sensitivity are essential to the operation of many solid-state qubits. These sensors can be realized using nanoelectronic devices such as single-electron transistors (SETs), quantum point-contacts (QPCs) and quantum dots, all of which are highly sensitive to their local electrostatic environment. The sensor can measure the state of a charge qubit directly [1, 2, 3, 4, 5] or, via spin-to-charge conversion, the state of a spin qubit [6, 7, 8, 9, 10, 11]. If the sensor bandwidth is large enough, the qubit state may be detected on a timescale shorter than its lifetime. Such single-shot detection is essential for implementing error correction in a quantum information processing architecture [12]. It can also be used to observe correlations between qubits [13, 14] and to measure the dynamics of the environment [15, 16]. Single-shot readout with a charge sensor has been realized for a variety of solid-state qubit candidates, such as semiconductor quantum dots in GaAs [17, 18, 19, 20, 6, 7, 21, 22, 23, 24, 25, 26], InAs nanowires [27], carbon nanotubes [28], graphene [29], silicon [30, 31, 32, 33, 34, 35], and superconducting charge qubits [36, 4]. In all cases the signal from the charge sensor takes one of two or more discrete levels, where each level corresponds to a different configuration of charge in the qubit.

Numerous sources of noise can pollute the charge sensor signal. Sources of white noise include shot noise and Johnson-Nyquist noise in warm parts of the measurement circuit. Amplifier noise, which is often dominant, may be approximated as white noise over limited frequency ranges [22]. An intrinsic source of $1/f$ noise is charge fluctuations in the vicinity of the sensor and the qubit [37]. Other potential noise sources include interference from AC power supplies, signal drift due to instability in the measurement electronics, and instability of the sensor itself. High fidelity readout of a qubit must therefore rely on the adoption of a filter or algorithm that can accurately identify real events in a noisy signal. Wavelet signal processing is one promising solution to this challenge.

Wavelet signal processing is a joint time-frequency analysis technique that is well-suited to identifying within a signal localised events that possess particular spectral characteristics. The technique has found numerous applications ranging from genetics to image compression and storage of fingerprint data [38, 39, 40]. One common use of wavelet analysis is the detection of sharp edges in images [41]. The problems of detecting edges in images and events in a charge-sensor signal are very similar: in both cases the signal of interest is a sharp step overlayed on a noisy background.

In this paper we assess the performance of a wavelet edge detection algorithm for identifying events in a single electron charge sensor signal using both experimental data and simulated signals. We compare wavelet detection to an approach based on a threshold signal value. Thresholding, and more sophisticated derivatives of this method, are commonly used to analyze charge sensor signals [4, 6, 7, 13, 14, 15, 20, 21, 24, 30, 31, 32, 34, 42]. Below we show that, while both techniques perform well at low noise, wavelet edge detection has much better tolerance to $1/f$ and low-frequency noise and is




Figure 1. Stages of wavelet edge detection. (a): first 200 ms of the input signal. The full signal has 2^{12} data points sampled at a rate of 2 kHz. It was generated by adding simulated white and $1/f$ noise to the “ideal” signal shown below the axis. (b): first 200 ms of the multi-scale wavelet transform of the signal (see main text for details). Local maxima and minima in the transform are tracked in the direction of decreasing wavelet scale s . Two example tracks are shown by green lines. A total ‘weight’ is calculated for each track (as described in the main text) and the results are shown in (c). Significant edges (open circles) are identified by tracks with a weight above 200 (the dashed line). Other edges (closed circles) are ignored. The position of an edge is the position of its track at the finest scale ($s = 2^2$). The edge direction (up or down) is given by the sign of the transform.

also more tolerant to white noise. The wavelet algorithm described below has previously been used to successfully analyze measurements of a Si/SiGe spin qubit with a QPC charge sensor [32]. The noise spectra used here to assess the detection techniques include examples that closely resemble those found in that work. Both techniques are tested over a range of noise profiles, and the results are applicable to any sensor where events are characterized by switches between discrete signal levels.

2. Method

Figure 1 outlines the stages of the wavelet edge detection process. The approach is both a formalization and generalization of Canny’s edge detection algorithm [43]. The detection is based upon a multi-scale real wavelet transform of the input signal. The transform W at a time t and wavelet scale s is found by convolving the signal $I(t)$ with the scaled mother wavelet ψ : $W(t, s) = \int_{-\infty}^{+\infty} I(t') \frac{1}{\sqrt{s}} \psi\left(\frac{t-t'}{s}\right) dt'$ [41]. By analogy with the Canny edge detection algorithm, we chose the mother wavelet ψ to be the first derivative of a Gaussian. The scale s sets the width of this wavelet and in Fig 1

and Fig 2 we express s as a number of data points. An upwards (downwards) step in the signal will produce a local maximum (minimum) in the wavelet transform, and significant edges result in large, local maxima and minima that persist across a broad range of wavelet scales [41].

Given a wavelet transform, the edge detection algorithm tracks local maxima and minima across wavelet scales to calculate a cumulative ‘weight’ for every maximum/minimum present at the smallest scale. The weight of a point on a track is defined as the square of the transform normalised by the median value of the square of the transform at that scale. Significant edges in the signal are identified by tracks with a total weight above a cutoff value, and this is the only parameter used to optimize the detection. Our implementation of the algorithm is realized in MATLAB and the wavelet transformations are performed using the WaveLab850 library [44].

Both experimental data and simulated charge sensor signals are used to investigate detection performance. The simulated signals are chosen to mimic typical measurements of Si/SiGe spin qubits using a QPC charge sensor [32, 34], and for easy comparison with typical data we present our results with a specific choice of time-scales on the abscissa; however, the results remain valid if either time and/or amplitude are scaled. The signals are generated by adding noise to an ideal charge sensor output. We choose a sampling rate of 2 kHz and a length of 2^{12} data points for all signals. The ideal output has two discrete levels at $\pm 0.5e$, where e is the electron charge. Switches between the levels are generated randomly according to a Poisson distribution with a characteristic rate of 100 Hz. A 1 kHz low-pass filter is applied to the ideal output, resulting in a signal rise-time that is common in experiments. After filtering, simulated white noise and $1/f$ noise are added to the signal.

The accuracy of the edge detection, when applied to a given signal, is measured by the F-score of the results $F = 2pr/(p + r)$, where the precision p is the fraction of detected events that correspond to real events, and the recall r is the fraction of events in the signal that were correctly detected [45]. When calculating p and r , we allow the maximum difference in time between real and detected events to be twice the time-constant of the low-pass filter (2 ms). $F = 1$ indicates perfect detection. Both false negatives and false positives reduce F .

3. Results and discussion

We investigate the performance of wavelet edge detection using both simulated signals and experimental data. By using artificial signals, the true position of every event in the signal is known and the detection accuracy can be calculated. This analysis is presented in section 3.1. In section 3.2 we investigate detection performance using measurements of a quantum point-contact charge sensor adjacent to a Si/SiGe double quantum dot. The measured signals are known to have characteristics that enable the performance of the edge detection to be estimated, without knowing the true positions of the charging events.




Figure 2. Comparative examples of wavelet and threshold edge detection. (a-c) Results for a signal with $1/f$ noise of RMS amplitude A_P . (d-e) Results for a signal with white noise of RMS amplitude A_W . (a) and (d) show the input signals. The optimum levels for threshold edge detection are shown by blue dashed lines. (b) and (e) show wavelet transforms of the signals. (c) and (f) show the location of events given by wavelet and threshold edge detection, below the ideal sensor output. Correctly detected events are marked by green arrows. Detected events that do not correspond to a real event are marked by red arrows.

3.1. Comparison between wavelet and threshold detection using simulated signals

We compare wavelet edge detection to a simple alternative thresholding algorithm in which an event is detected when the signal crosses a threshold. We optimize both detection techniques for each signal by varying a single parameter to maximize F . In the case of threshold edge detection, this parameter is the threshold level. In the case of wavelet edge detection, the parameter is the final cutoff that determines whether a track in the transform has a large enough weight to be accepted as an event.

Figure 2 shows the results of applying wavelet edge detection and the thresholding technique to two example signals. The signal in Fig. 2(a) was generated with $1/f$ noise of RMS amplitude of $A_P = 0.25e/\sqrt{\text{Hz}}$ at 1 Hz and zero white noise. Wavelet edge detection performs well in this situation, achieving a score of $F = 0.99$ over the full length of the signal. Threshold edge detection is significantly affected by $1/f$ noise, reaching only $F = 0.79$. This low score is easily understood: large, low-frequency




Figure 3. Averaged F-score (see definition in main text) for wavelet edge detection (b),(d) and threshold edge detection (a),(c) as a function of white noise amplitude A_W and $1/f$ noise amplitude A_P . (c) and (d) show detail within the black dashed boxes in (a) and (b) respectively. A score of $F = 1$ corresponds to perfect detection. A score of $F = 2/3$ corresponds, for example, to half the edges being detected with no false positives, or all edges being detected but with an equal number of false positives. The wavelet edge detection performs as well as threshold edge detection or better at all noise levels, and is significantly more robust against $1/f$ noise.

components of the noise cause the signal to drift with respect to the threshold level, resulting in false positives and missed events. In contrast, the same noise does little to distort the sharp edges in the data and these are easily identified in the wavelet transform.

The signal in Fig. 2(d) was generated with zero $1/f$ noise and white noise of RMS amplitude $A_W = 0.008e/\sqrt{\text{Hz}}$ (a peak-to-peak noise amplitude of $1.01e$ in this bandwidth). Over the full length of the signal, the wavelet edge detection has a score of $F = 0.88$ and the threshold edge detection has $F = 0.40$. The low score for thresholding is largely due to false positive detections, which is unsurprising given the low signal-to-noise ratio. The wavelet edge detection is still able to identify events because edges produce a distinct profile in the wavelet transform that is not reproduced by white noise: a local maximum that extends over a large range of scales.

The two detection techniques were compared over a range of noise amplitudes. Fig. 3(a-d) shows F-scores for each value of (A_W, A_P) , which were found by averaging over 10 signals that were generated by adding random noise to the same ideal sensor output. Each signal is 2^{12} data points (≈ 2 s) long, and each plot in Fig. 3 is the result of

analyzing $\sim 10,000$ signals. The results show that the wavelet edge detection matches or outperforms thresholding at all noise levels. From Fig. 3(c) and (d), it is clear that the wavelet approach is significantly more tolerant to $1/f$ noise in the signal and, by extension, low-frequency noise in general.

There are several improvements that may be made to the wavelet edge detection algorithm. Typically, the slowest part of the analysis is tracking edges in the wavelet transform. The analysis time can be significantly reduced by choosing only to track maxima/minima in the transform once they exceed a certain weight. This ensures that large numbers of short, weak tracks, such as those produced by white noise, are discarded. The rate of false positives may be reduced by including a prior knowledge of the detector sensitivity: edges can be rejected if the signal level does not change by the expected value in the vicinity of the edge. It should be noted that there are also many ways to improve the thresholding technique presented here, although at the cost of adding additional parameters that must be tuned; for example, local averaging to compensate for low-frequency signal drift [17], or using a Schmitt trigger to select events [42]. In general, these filters will make threshold edge detection more similar to the wavelet approach, which can be thought of as a filter with an excellent selectivity to sharp edges.

3.2. Comparison between wavelet and threshold detection using experimentally measured signals

To demonstrate the benefits of wavelet edge detection we apply the technique to experimental data obtained from a point-contact charge sensor adjacent to a Si/SiGe double quantum dot. Fig. 4 shows an SEM image of the device. The charge sensor signal I_{QPC} has two preferred values, corresponding to an additional electron being in either the left or right quantum dot. Further details of the device and the measurement setup can be found in Prance et al. [34]. Importantly for this work, events in the charge sensor signal are known to occur at certain times, as discussed below, which makes it possible to independently estimate the accuracy of the edge detection. Based on this *a priori* knowledge of the signal, we can compare the accuracy of the wavelet and thresholding algorithms using real data.

An example signal from the Si/SiGe point-contact charge sensor is shown in Fig. 4. In contrast to the data used in section 3.1, the times at which edges occur in this signal are not random. During the measurement the double dot was driven by a 300 Hz square wave. Any charging events in the signal are aligned to this pulse with a high probability and there are only two expected outcomes within each pulse period: either no edges or an up-down pair of edges will be present. All other outcomes are known to be unlikely due to the behaviour of the double quantum dot [34].

We use wavelet and threshold edge detection to find the number of edges during each pulse period. We estimate the error of each edge detection method to be the percentage of pulse periods that are found to contain an unexpected number of edges




Figure 4. Example charge sensor data from a quantum point-contact adjacent to a Si/SiGe double quantum dot [34]. (a) Measured current I_{QPC} through the charge sensor as the quantum dots are driven by a 300 Hz square wave. Either two or zero charging events are expected during each ≈ 3.3 ms pulse period. The periods are separated by vertical grey lines, and the observed number of events in each period are shown above the data. Pairs of events are visible in the signal as peaks (one upwards transition and one downwards). (b) Positions and directions of edges detected in the signal by the thresholding and wavelet algorithms. Wavelet detection produced no errors in this section of signal. Thresholding detection produced several errors, indicated by red arrows. In this section of the signal, the majority of these errors are multiple detections of the same edge due to high frequency signal noise.

(i.e. not zero or two). Table 1 shows the optimised errors using wavelet and thresholding edge detection as described in sections 2 and 3.1. The values in Table 1 are based on an analysis of 120 signals, each consisting of 10^5 data points (a total of $\approx 72,000$ pulse periods) sampled at a rate of 50 kHz. For this dataset the wavelet detection is found to significantly outperform the simple thresholding algorithm. (The wavelet error is $\approx 2\%$ while the thresholding error is $\approx 50\%$).

The performance of the thresholding algorithm can be improved at the cost of reduced bandwidth, extra free parameters, and the need for additional post-processing. If the threshold is defined relative to the average value of the signal, the effect of low frequency signal drift can be reduced. This drift correction decreases the error marginally to $\approx 49\%$. A greater improvement is found by applying a low-pass filter. When the signal bandwidth is reduced to 10 kHz, the error is reduced to $\approx 14\%$. Combining this low-pass filter with drift correction gives an error of $\approx 13\%$. The low pass filter has negligible effect on the wavelet edge detection, which already filters the signal at multiple scales as part of its operation. While it is possible to achieve reasonable results using thresholding on this data set, the resulting algorithm includes one additional free

Edge detection method	Unexpected events
Wavelet	2.26%
Wavelet with 10 kHz filter	2.24%
Simple threshold	50.2%
Threshold with drift correction (d.c.)	48.7%
Threshold with 10 kHz filter	13.7%
Threshold with 10 kHz filter and d.c.	12.5%

Table 1. Estimate of the characteristic error for several edge detection algorithms when applied to experimental charge sensor data. The error is estimated by comparing the detected edges to the expected characteristics of the signal (see Fig. 4). For each result the free parameters of the corresponding algorithm were optimised to minimise the error.

parameter, the filter cut-off frequency, and needs knowledge of the mean signal level, which requires averaging over long times. Furthermore, while the 10 kHz filter reduces the error in this particular data set, it limits the detection bandwidth in general. By contrast, the wavelet edge detection performs well with just one free parameter and no additional filtering.

4. Conclusions

We have compared the performance of a wavelet edge detection algorithm to a thresholding methods using simulated signals that mimic the output of a QPC charge sensor and measured signals from a Si/SiGe double quantum dot device. When only a single free parameter is allowed in the detection algorithms, wavelet edge detection is found to outperform thresholding over all noise profiles in the simulated signals and is particularly robust against $1/f$ noise. Wavelet detection is also found to perform well on real data, even when compared to a more complex thresholding algorithm with additional free parameters. A significant practical advantage is that, once optimized, the wavelet edge detection can be used in the presence of low-frequency noise and signal drift without adjustment. This means that the analysis may be performed while the experiment is running, which is crucial for certain qubit applications such as error correction. Furthermore, the excellent sensitivity of wavelet edge detection can be used to increase the fidelity and potentially improve the bandwidth of qubit readout based on a single electron charge sensor.

Acknowledgments

This research was supported in part by the U.S. Army Research Office (W911NF-08-1-0482, W911NF-12-1-0607). Development and maintenance of the growth facilities used for fabricating samples is supported by DOE (DE-FG02-03ER46028). This research utilized NSF-supported shared facilities at the University of Wisconsin-Madison.

References

- [1] Field M, Smith C G, Pepper M, Ritchie D A, Frost J E F, Jones G A C and Hasko D G 1993 *Phys. Rev. Lett.* **70** 1311–1314
- [2] Elzerman J M, Hanson R, Greidanus J S, Willems van Beveren L H, De Franceschi S, Vandersypen L M K, Tarucha S and Kouwenhoven L P 2003 *Phys. Rev. B* **67** 161308
- [3] Petta J R, Johnson A C, Marcus C M, Hanson M P and Gossard A C 2004 *Phys. Rev. Lett.* **93** 186802
- [4] Astafiev O, Pashkin Y A, Yamamoto T, Nakamura Y and Tsai J S 2004 *Phys. Rev. B* **69**(18) 180507
- [5] Shi Z, Simmons C B, Ward D R, Prance J R, Mohr R T, Koh T S, Gamble J K, Wu X, Savage D E, Lagally M G, Friesen M, Coppersmith S N and Eriksson M A 2013 *Phys. Rev. B* **88**(7) 075416
- [6] Elzerman J M, Hanson R, Willems van Beveren L H, Witkamp B, Vandersypen L M K and Kouwenhoven L P 2004 *Nature* **430** 431–435
- [7] Hanson R, van Beveren L H W, Vink I T, Elzerman J M, Naber W J M, Koppens F H L, Kouwenhoven L P and Vandersypen L M K 2005 *Phys. Rev. Lett.* **94** 196802
- [8] Johnson A C, Petta J R, Taylor J M, Yacoby A, Lukin M D, Marcus C M, Hanson M P and Gossard A C 2005 *Nature* **435** 925–928
- [9] Petta J, Johnson A, Taylor J, Laird E, Yacoby A, Lukin M, Marcus C, Hanson M and Gossard A 2005 *Science* **309** 2180–2184
- [10] Shi Z, Simmons C B, Prance J R, Gamble J K, Koh T S, Shim Y P, Hu X, Savage D E, Lagally M G, Eriksson M A, Friesen M and Coppersmith S N 2012 *Phys. Rev. Lett.* **108**(14) 140503
- [11] Kim D, Shi Z, Simmons C B, Ward D R, Prance J R, Koh T S, Gamble J K, Savage D E, Lagally M G, Friesen M, Coppersmith S N and Eriksson M A 2014 *Nature* **511** 70–74
- [12] Ladd T D, Jelezko F, Laflamme R, Nakamura Y, Monroe C and O’Brien J L 2010 *Nature* **464** 45–53
- [13] Nowack K C, Shafiei M, Laforest M, Prawiroatmodjo G E D K, Schreiber L R, Reichl C, Wegscheider W and Vandersypen L M K 2011 *Science* **333** 1269–1272
- [14] Shulman M D, Dial O E, Harvey S P, Bluhm H, Umansky V and Yacoby A 2012 *Science* **336** 202–205
- [15] Barthel C, Medford J, Bluhm H, Yacoby A, Marcus C M, Hanson M P and Gossard A C 2012 *Phys. Rev. B* **85**(3) 035306
- [16] Fink T and Bluhm H 2013 *Phys. Rev. Lett.* **110**(1) 010403
- [17] Lu W, Ji Z, Pfeiffer L, West K W and Rimberg A J 2003 *Nature* **423** 422–425 ISSN 0028-0836
- [18] Fujisawa T, Hayashi T, Hirayama Y, Cheong H D and Jeong Y H 2004 *Appl. Phys. Lett.* **84** 2343–2345
- [19] Vandersypen L M K, Elzerman J M, Schouten R N, van Beveren L H W, Hanson R and Kouwenhoven L P 2004 *Appl. Phys. Lett.* **85** 4394–4396
- [20] Schleser R, Ruh E, Ihn T, Ensslin K, Driscoll D C and Gossard A C 2004 *Appl. Phys. Lett.* **85** 2005–2007
- [21] Meunier T, Vink I T, Willems van Beveren L H, Koppens F H L, Tranitz H P, Wegscheider W, Kouwenhoven L P and Vandersypen L M K 2006 *Phys. Rev. B* **74**(19) 195303
- [22] Gustavsson S, Leturcq R, Simovic B, Schleser R, Ihn T, Studerus P, Ensslin K, Driscoll D C and Gossard A C 2006 *Phys. Rev. Lett.* **96** 076605 (pages 4)
- [23] Reilly D J, Marcus C M, Hanson M P and Gossard A C 2007 *Appl. Phys. Lett.* **91** 162101 (pages 3)
- [24] Barthel C, Reilly D J, Marcus C M, Hanson M P and Gossard A C 2009 *Phys. Rev. Lett.* **103**(16) 160503
- [25] Gustavsson S, Leturcq R, Studer M, Shorubalko I, Ihn T, Ensslin K, Driscoll D and Gossard A 2009 *Surf. Sci. Rep.* **64** 191 – 232 ISSN 0167-5729
- [26] Mason J, Gaudreau L, Studenikin S, Kam A, Djurkovic B, Sachrajda A and Kycia J 2010 *Physica*

E **42** 813 – 816 ISSN 1386-9477

- [27] Gustavsson S, Shorubalko I, Leturcq R, Schön S and Ensslin K 2008 *Appl. Phys. Lett.* **92** 152101 (pages 3)
- [28] Gotz G, Steele G A, Vos W J and Kouwenhoven L P 2008 *Nano Letters* **8** 4039–4042 pMID: 18928322
- [29] Güttinger J, Seif J, Stampfer C, Capelli A, Ensslin K and Ihn T 2011 *Phys. Rev. B* **83** 165445
- [30] Morello A, Pla J J, Zwanenburg F A, Chan K W, Tan K Y, Huebl H, Möttönen M, Nugroho C D, Yang C, van Donkelaar J A, Alves A D C, Jamieson D N, Escott C C, Hollenberg L C L, Clark R G and Dzurak A S 2010 *Nature* **467** 687–691
- [31] Thalakulam M, Simmons C B, Van Bael B J, Rosemeyer B M, Savage D E, Lagally M G, Friesen M, Coppersmith S N and Eriksson M A 2011 *Phys. Rev. B* **84**(4) 045307
- [32] Simmons C B, Prance J R, Van Bael B J, Koh T S, Shi Z, Savage D E, Lagally M G, Joynt R, Friesen M, Coppersmith S N and Eriksson M A 2011 *Phys. Rev. Lett.* **106** 156804
- [33] Mahapatra S, Buch H and Simmons M Y 2011 *Nano Letters* **11** 4376–4381
- [34] Prance J R, Shi Z, Simmons C B, Savage D E, Lagally M G, Schreiber L R, Vandersypen L M K, Friesen M, Joynt R, Coppersmith S N and Eriksson M A 2012 *Phys. Rev. Lett.* **108**(4) 046808
- [35] Yuan M, Yang Z, Savage D E, Lagally M G, Eriksson M A and Rimberg A J 2012 *Appl. Phys. Lett.* **101** 142103 (pages 3)
- [36] Schoelkopf R J, Wahlgren P, Kozhevnikov A A, Delsing P and Prober D E 1998 *Science* **280** 1238–1242
- [37] Jung S W, Fujisawa T, Hirayama Y and Jeong Y H 2004 *Appl. Phys. Lett.* **85** 768–770
- [38] The ENCODE Project Consortium 2007 *Nature* **447** 799–816
- [39] Joint Photographic Experts Group JPEG 2000 standard URL <http://www.jpeg.org/jpeg2000>
- [40] Bradley J N, Brislaw C M and Hopper T 1993 *Proc. SPIE* **1961** 293–304
- [41] Mallat S 2009 *A Wavelet Tour of Signal Processing* 3rd ed (Elsevier Academic Press)
- [42] Lambert N J, Edwards M, Esmail A A, Pollock F A, Barrett S D, Lovett B W and Ferguson A J 2013 *ArXiv e-prints* (Preprint 1304.5117)
- [43] Canny J 1986 *Pattern Analysis and Machine Intelligence, IEEE Transactions on PAMI* **8** 679–698 ISSN 0162-8828
- [44] The function ‘RWT’ in WaveLab850 is used to calculate the wavelet transform. The function ‘MM_RWT’ is used to find the location of local maxima in the modulus of the transform. URL <http://www-stat.stanford.edu/~wavelab>
- [45] van Rijsbergen C J 1979 *Information Retrieval* 2nd ed (Butterworths)